

MULTILEVEL ENSEMBLE TRANSFORM PARTICLE FILTERING

A. GREGORY [†], C. J. COTTER [†] AND S. REICH [‡]

Abstract. This paper extends the Multilevel Monte Carlo variance reduction technique to nonlinear filtering. In particular, Multilevel Monte Carlo is applied to a certain variant of the particle filter, the Ensemble Transform Particle Filter. A key aspect is the use of optimal transport methods to re-establish correlation between coarse and fine ensembles after resampling; this controls the variance of the estimator. Numerical examples present a proof of concept of the effectiveness of the proposed method, demonstrating significant computational cost reductions (relative to the single-level ETPF counterpart) in the propagation of ensembles.

Key words. Multilevel Monte Carlo, sequential data assimilation, optimal transport

AMS subject classifications. 65C05, 62M20, 93E11, 93B40, 90C05

1. Introduction. Data assimilation is the process of incorporating observed data into model forecasts. In data assimilation, one is interested in computing statistics $\mathbb{E}_\eta[X]$ of solutions X to random dynamical systems with respect to a posterior measure (η) given partial observations of the system. In particle filtering [7, 4], this is done by using an empirical ensemble representing the posterior distribution η at any one time. The propagation in time of the members (particles) of this ensemble can be computationally expensive, especially in high dimensional systems.

Recently, the Multilevel Monte Carlo (MLMC) method has been developed for achieving significant cost reductions in Monte Carlo simulations [9]. It has been applied to areas such as Markov Chain Monte Carlo [13] and quasi-Monte Carlo [10] to return computational cost reductions from existing techniques. It has also been applied to uncertainty quantification within PDEs [6]. The idea is to consider a hierarchy of discretized models, balancing numerical error in cheap/coarse models against Monte Carlo variance in expensive/fine models. It is desirable to adapt MLMC to sequential Monte Carlo methods such as particle filters, and some first steps have been taken in this direction. Firstly, [11] have developed a multilevel Ensemble Kalman Filter (EnKF), using MLMC estimators to calculate the mean and covariance of the posterior, in the case where the underlying distributions are Gaussian and the model is linear. However for non-Gaussian distributions and nonlinear models, the EnKF is biased. The method does however converge to a “mean-field limit” [14]. Secondly, [3] proposed a multilevel sequential Monte Carlo method for Bayesian inference problems to give significant computational cost reductions from standard techniques. Our goal in this paper is to take a step further by applying MLMC to nonlinear filtering problems.

[†]Department of Mathematics, Imperial College London, Exhibition Road, London SW7 2AZ, UK (a.gregory14@imperial.ac.uk). Alastair Gregory was supported by the Science and Solutions to a Changing Planet DTP and the Natural Environmental Research Council.

[‡]Institut für Mathematik, Universität Potsdam, Am Neuen Palais 10, D-14469 Potsdam, Germany and Department of Mathematics and Statistics, University of Reading, Whiteknights, Reading RG6 6AX, UK.

In general, MLMC works by computing statistics from pairs of coarser and finer correlated ensembles. For Monte Carlo simulation of SDEs, this correlation is achieved by using the same initial conditions and Brownian paths for each coarse/fine pair of ensemble members. The key challenge in applying MLMC to particle filtering is in maintaining this correlation after resampling. [8] suggested that correlating coarse and fine ensembles could be achieved by minimising the Wasserstein distance between the two ensembles. This can be formulated as an optimal transportation problem [20].

In this paper, we adapt the MLMC framework to the Ensemble Transform Particle Filter (ETPF) [19]. ETPF is an efficient and effective nonlinear filter that uses optimal transportation transformations [22] instead of random resampling. In our Multilevel ETPF (MLETPF) the coupling between coarse and fine ensembles is also maintained using optimal transportation. The sole aim of introducing MLMC to the ETPF is to reduce the computational cost of the propagation of particles. This is only a benefit if the computational cost dominates the optimal transportation transformation cost; whilst direct solvers for optimal transportation problems with one-dimensional state space scale as $O(N \log(N))$, solvers for problems with more than one dimension scale as $O(N^3 \log(N))$ with the ensemble size. To address this, a technique commonly used in the aforementioned EnKF known as localisation can be used to reduce this optimal transportation cost significantly [5]. Our proposed MLETPF can return significant reductions in the overall computational cost of ETPF where the particle propagation cost dominates. It will also return significant reductions in cases where optimal transportation computational cost dominates, if the localised ETPF is used.

This paper proceeds as follows. Section 2 provides a background of the MLMC method, §3 describes the basic particle filtering framework together with the ETPF scheme. Then, the proposed Multilevel ETPF (MLETPF) method is presented in §4, along with numerical examples to demonstrate the effectiveness of the method. Finally, §5 provides a summary and outlook.

2. The Multilevel Monte Carlo Method. The Multilevel Monte Carlo estimator can be viewed as a variance reduction technique for a standard Monte Carlo estimator. Suppose one wishes to compute an approximation of $\mathbb{E}[X_L]$, where X_L is a numerical approximation of a random variable X (with discretization accuracy parameter¹ $h_L \propto M^{-L}$). The Multilevel Monte Carlo (MLMC) method introduced in [9] considered the case where X is the solution to a stochastic differential equation (SDE) at time $T > 0$; the discretized solutions X_L are obtained from a given numerical method with stepsize h_L . This paper will instead consider X to be a solution, at time $T > 0$, to a general random dynamical system, with stochastic forcing and/or random initial conditions (drawn from a distribution π^0). In the simplest case let X_L^i , $i = 1, \dots, N$, be $N \geq 1$ i.i.d. samples of X_L . The standard, unbiased, Monte Carlo

¹Such as the time stepsize.

estimator for $\mathbb{E}[X_L]$ is then

$$(2.1) \quad \bar{X}_L^{MC} = \frac{1}{N} \sum_{i=1}^N X_L^i.$$

Using a telescoping sum of expectations,

$$(2.2) \quad \mathbb{E}[X_L] = \mathbb{E}[X_0] + \sum_{l=1}^L \mathbb{E}[X_l] - \mathbb{E}[X_{l-1}],$$

one can define the MLMC approximation to $\mathbb{E}[X_L]$ as a sum of independent Monte Carlo estimators, $\bar{X}_L = \sum_{l=0}^L \hat{X}_l$, where

$$(2.3) \quad \hat{X}_l = \begin{cases} \sum_{i=1}^{N_0} \frac{1}{N_0} X_0^i, & l = 0, \\ \sum_{i=1}^{N_l} \frac{1}{N_l} (X_l^i - X_{l-1}^i), & l > 0, \end{cases}$$

leading to

$$(2.4) \quad \bar{X}_L = \frac{1}{N_0} \sum_{i=1}^{N_0} X_0^i + \sum_{l=1}^L \left(\frac{1}{N_l} \sum_{i=1}^{N_l} (X_l^i - X_{l-1}^i) \right).$$

Here, N_l , $l = 0, \dots, L$, are Monte Carlo sample sizes, of i.i.d. draws from X_l, X_{l-1} respectively, for each of the $L + 1$ estimators. We have

$$(2.5) \quad \mathbb{E}[\hat{X}_l] = \begin{cases} \mathbb{E}[X_0], & l = 0, \\ \mathbb{E}[X_l] - \mathbb{E}[X_{l-1}], & l > 0, \end{cases}$$

and hence the MLMC estimator is an unbiased approximation of $\mathbb{E}[X_L]$. We will call the estimators, \hat{X}_l , $l > 0$, ‘multilevel difference estimators’. The important thing to note here is that the fine (level l) and coarse (level $l - 1$) samples in each difference estimator must be positively correlated for each i in each of the $L + 1$ multilevel difference estimators. This can be achieved by using the same random system input (e.g. initial conditions/stochastic forcing) for each i on both levels. On the other hand, the samples in different difference estimators must be uncorrelated.

For fixed T , the discretization bias (away from $\mathbb{E}[X]$) of the overall estimator is $O(h_L^\alpha)$ [9], where α is the global discretization bias (i.e. $|\mathbb{E}[X_L] - \mathbb{E}[X]|$) of the numerical method used to simulate X_l , $l \geq 0$. One notes from [9] that,

$$(2.6) \quad |\mathbb{E}[X_l] - \mathbb{E}[X_{l-1}]| \leq (M - 1)ch_l^\alpha$$

where c is a positive constant, and that $(M - 1)^{-1}|\mathbb{E}[X_L] - \mathbb{E}[X_{L-1}]|$ can be used as an estimate for the overall discretization bias, $|\bar{X}_L - \mathbb{E}[X]|$. As each estimator in (2.4) is independent of one another, the overall variance is given by the sum of the variances of each individual estimator. Given that there is a positive correlation between X_l and X_{l-1} , one can expect then that the sample variance of $X_l - X_{l-1}$, denoted by

$$(2.7) \quad V_l = \mathbb{V}[X_l - X_{l-1}] = \mathbb{V}[X_l] + \mathbb{V}[X_{l-1}] - 2\text{Cov}[X_l, X_{l-1}],$$

decays at a rate proportional to l , so that $\mathbb{V}_l = O(h_l^\beta)$, $\beta > 0$. The covariance in the last term in (2.6) is taken over the joint probability distribution of X_l and X_{l-1} . One can then trade off variance in fine/expensive estimators against discretization error in coarse/cheap estimators with lower variance by setting a decreasing sequence of Monte Carlo estimator sample sizes, $N_0 > N_1 \dots > N_L$. The overall computational cost of the MLMC estimator is

$$(2.8) \quad C_{ML} = \sum_{l=0}^L h_l^{-\gamma} N_l T,$$

where, $h_l^{-\gamma}$ defines the computational cost of propagating one single sample (of X_l) through a discretized system. On the other hand, the cost of the standard Monte Carlo estimator in (2.1) is

$$(2.9) \quad C_{MC} = h_L^{-\gamma} N T.$$

One can choose the continuous variables N_l , $l = 0, \dots, L$, at a rate that minimises the variance of the MLMC estimator for a fixed computational cost, $N_l \propto \sqrt{V_l h_l^\gamma}$. In particular, by following a formula given in [9, 6], one can find optimal values of N_l , as well as the finest level L , and in doing so achieve a computational cost reduction relative to the standard Monte Carlo counterpart (2.1), with the same bound Mean Square error. Giles [9] proved the following result.

THEOREM 2.1. *If the Mean Square Error of \bar{X}_L is bounded by $O(\varepsilon^2)$, one can optimally choose L and N_l to allow the computational cost of the MLMC estimator to be bounded by,*

$$(2.10) \quad C_{ML} \leq \begin{cases} c_1 \varepsilon^{-2}, & \gamma < \beta \\ c_2 \varepsilon^{-2} \log(\varepsilon)^2, & \gamma = \beta \\ c_3 \varepsilon^{-2 - \frac{(\gamma-\beta)}{\alpha}}, & \gamma > \beta \end{cases}$$

where $c_1, c_2, c_3, \gamma, \beta$ are positive constants and $\alpha \geq \frac{1}{2} \min(\gamma, \beta)$.

For \bar{X}_L^{MC} to have a Mean Square Error of $O(\varepsilon^2)$, a sample size N of $O(\varepsilon^{-2})$ is required, as well as a discretization bias given by $h_L = O(\varepsilon^{\frac{1}{\alpha}})$. Thus any of the computational costs in Theorem 2.1 are less than C_{MC} ($O(\varepsilon^{-2 - \frac{1}{\alpha}})$). The MLMC approach in principle is very simple to implement and can be very effective as long as one can satisfy the two constraints, $\alpha \geq \frac{1}{2} \min(\gamma, \beta)$ and $\beta > 0$. More detail on this method can be found in a generalised explanation, and related theorems, in [6].

3. Particle Filtering. This section will outline the standard particle filtering methodology. In this context, one is interested in computing statistics of a random process X_t , conditioned on observations of a single realisation of X , denoted X' , and referred to as the reference solution. The observations are random variables of the form

$$(3.1) \quad Y_{t_k} = H(X'_{t_k}) + \phi$$

at times $t \in (t_1, \dots, t_{N_y})$, where $H : \mathbb{R}^d \rightarrow \mathbb{R}^y$ is an observation operator, and ϕ is a random variable representing measurement error. For simplicity, we choose $\phi \sim N(0, R)$, R is a $y \times y$ covariance matrix, and we will take H to be the identity (so $y=d$). We define X_{L,t_k} to be a numerical discretization of X_{t_k} with discretization accuracy parameter h_L . Our aim here is to sequentially approximate $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, the expectation of X_{L,t_k} with respect to the measure η_{L,t_k} , where η_{L,t_k} is the posterior of X_{L,t_k} given the observations Y_{t_1, \dots, t_k} . Let $p(y|x)$ be the likelihood function of y given x and $q(x)$ be the prior of x . Then, for any $k \in [1, N_y]$, using Importance Sampling [7], one can draw N i.i.d samples from the empirical approximation of the prior of X_{L,t_k} ,

$$(3.2) \quad \hat{q}(X_{L,t_k}) = \sum_{i=1}^N \tilde{w}_L(X_{L,t_{k-1}}^i) \delta(X_{L,t_k} - X_{L,t_k}^i),$$

and denote the normalised importance weights of sample $i \in [1, N]$ to be

$$(3.3) \quad \tilde{w}_L(X_{L,t_k}^i) = \frac{w_L(X_{L,t_k}^i)}{\sum_{j=1}^N w_L(X_{L,t_k}^j)},$$

where

$$(3.4) \quad w_L(X_{L,t_k}^i) = p(Y_{t_k} | X_{L,t_k}^i) \tilde{w}_L(X_{L,t_{k-1}}^i).$$

Thus, the filter weights are defined iteratively, starting from $\tilde{w}_L(X_{L,t_0}^i) = \frac{1}{N}$. As the observations are given by a Gaussian distribution, the likelihood is

$$(3.5) \quad p(Y_{t_k} | X_{L,t_k}^i) = \frac{1}{\sqrt{2\pi|R|^{y/2}}} e^{-\frac{1}{2} (H(X_{L,t_k}^i) - Y_{t_k})^T R^{-1} (H(X_{L,t_k}^i) - Y_{t_k})}.$$

Finally, an estimator for the expectation of X_{L,t_k} with respect to the posterior η_{L,t_k} is

$$(3.6) \quad \bar{X}_{L,t_k} = \sum_{i=1}^N \tilde{w}_L(X_{L,t_k}^i) X_{L,t_k}^i.$$

This estimator, despite being biased by $O(N^{-1})$ due to the normalised importance weights in a single importance sampling update, is consistent with $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$. This is to say, as $N \rightarrow \infty$, the estimator converges in probability to $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$.

Typically, importance weights become degenerate as k increases [4]. In this case, it is necessary to duplicate higher weighted particles whilst removing lower weighted particles; this is known as resampling. Resampling resembles an unbiased transformation from the weighted ensemble, $\{X_{L,t_k}^i, \tilde{w}_L(X_{L,t_k}^i)\}_{i=1, \dots, N}$ to an evenly weighted ensemble of resampled particles $\{\tilde{X}_{L,t_k}^i\}_{i=1, \dots, N}$. The scheme outlined above is known as the Sequential Importance Resampling (SIR) method. For more information on this turn to [7]. The Ensemble Transform Particle Filter, the subject of this paper, uses Optimal Transportation [22] to implement this transformation, which we describe next.

3.1. Ensemble Transform Particle Filters. ETPFs are a variant of linear ensemble transform filters (LETF) [19]. They present an alternative to the resampling step that takes place in the standard SIR methodology, replacing it with a linear transformation. The goal is to obtain a transformed set of evenly weighted particles, $\{\tilde{X}_{L,t_k}^i\}_{i=1,\dots,N}$, from the weighted set of particles $\{X_{L,t_k}^i\}_{i=1,\dots,N}$, with importance weights $\{\tilde{w}_L(X_{L,t_k}^i)\}_{i=1,\dots,N}$, defining an empirical approximation to the posterior distribution η_{L,t_k} . This can be done with the following linear transformation,

$$(3.7) \quad \tilde{X}_{L,t_k}^j = \sum_{i=1}^N P_{i,j} X_{L,t_k}^i$$

for $i = 1, \dots, N$ and $j = 1, \dots, N$ with non-zero entries for $P_{i,j}$. Here, $\sum_{i=1}^N P_{i,j} = 1$. Let Z_{L,t_k} denote the discrete random variable with samples X_{L,t_k}^i and associated probability vector $\tilde{w}_L(X_{L,t_k}^i)$, $i = 1, \dots, N$. Then take \tilde{Z}_{L,t_k} to be the discrete random variable with samples X_{L,t_k}^i , $i = 1, \dots, N$ all with equal probability. For the ETPF, one creates a coupling between Z_{L,t_k} and \tilde{Z}_{L,t_k} , denoted by the matrix $T_{i,j}$, size $N \times N$, with non-negative entries. The coupling defines the linear transformation matrix in (3.7) as $P_{i,j} = NT_{i,j}$. This coupling can be found by solving a linear transport problem by minimising the expected Euclidean distance between Z_{L,t_k} and \tilde{Z}_{L,t_k} , subject to the constraints

$$(3.8) \quad \sum_{i=1}^N T_{i,j} = \frac{1}{N}, \quad \sum_{j=1}^N T_{i,j} = \tilde{w}_L(X_{L,t_k}^i).$$

This is in fact equivalent to maximising the covariance between the two ensembles, since

$$(3.9) \quad \begin{aligned} \mathbb{E}_{Z_{L,t_k}, \tilde{Z}_{L,t_k}} [\|z_{L,t_k} - \tilde{z}_{L,t_k}\|^2] &= \mathbb{E}_{Z_{L,t_k}} [\|z_{L,t_k}\|^2] + \mathbb{E}_{\tilde{Z}_{L,t_k}} [\|\tilde{z}_{L,t_k}\|^2] \dots \\ &\dots - 2\mathbb{E}_{Z_{L,t_k}} [z_{L,t_k}]^T \mathbb{E}_{\tilde{Z}_{L,t_k}} [\tilde{z}_{L,t_k}] - 2\text{Tr}(\text{Cov}_{Z_{L,t_k}, \tilde{Z}_{L,t_k}} [z_{L,t_k}, \tilde{z}_{L,t_k}]). \end{aligned}$$

In a univariate case, we define an optimal coupling matrix, $T_{i,j}$, as one which minimises the cost function,

$$(3.10) \quad \sum_{i=1}^N \sum_{j=1}^N T_{i,j} (X_{L,t_k}^i - X_{L,t_k}^j)^2$$

Theoretical analysis of the above transformation is given in [19]. Once the transformed particles in (3.7) are found, the posterior mean is now estimated by,

$$(3.11) \quad \bar{X}_{L,t_k} = \frac{1}{N} \sum_{j=1}^N \tilde{X}_{L,t_k}^j.$$

It is important to note that this linear transformation, which is deterministic, will give the same estimator as in (3.6), and thus does not add considerable extra variance to

the estimator from random resampling. This is a consistent estimator for the previous posterior mean estimator $\left(\sum_{i=1}^N \tilde{w}_L(X_{L,t_k}^i)X_{L,t_k}^i\right)$, since

$$(3.12) \quad \bar{X}_{L,t_k} = \frac{1}{N} \sum_{j=1}^N \tilde{X}_{L,t_k}^j = \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^N T_{i,j} N X_{L,t_k}^i = \sum_{j=1}^N \sum_{i=1}^N T_{i,j} X_{L,t_k}^i = \sum_{i=1}^N \tilde{w}_L(X_{L,t_k}^i) X_{L,t_k}^i$$

In the univariate case, the matrix $T_{i,j}$ can easily be found by an $O(N \log(N))$ algorithm in [20]. This will become an important observation when discussing localisation in the next section. The above constraints lead to a maximum of $2N - 1$ non-zero elements in $T_{i,j}$, leading to a very sparse matrix calculation, and thus the ensemble transformation process can be achieved in a $O(N)$ computational cost. The $O(N \log(N))$ computational cost comes from the fact that one has to sort the univariate particles prior to the algorithm. In our numerical experiments at the end of this paper, this sorting was a negligible part of the ensemble transform computational cost. This allows one to be able to carry the ensemble transform out on every assimilation step without the computational expense of this being of a higher order of magnitude than the propagation of the particles in between assimilation steps, but more analysis will cover this observation in the next section.

In the multivariate case, the same linear transport problem prevails, however one is required to minimise the cost function,

$$(3.13) \quad \sum_{i=1}^N \sum_{j=1}^N T_{i,j} \left\| X_{L,t_k}^i - \tilde{X}_{L,t_k}^j \right\|^2,$$

whose minimum defines the Wasserstein distance between Z_{L,t_k} and \tilde{Z}_{L,t_k} . This can be solved in a $O(N^3 \log(N))$ computational cost, using algorithms such as the FastEMD algorithm [16]. However, this means that with many systems, this optimal transportation computational cost will dominate over model costs of the system and thus the scheme is not efficient. The model costs of the particle filter are defined to be the computational cost needed to propagate the N particles through the system in between assimilation steps and is given in (2.9) (determined by a constant γ). Thankfully, a technique called *localisation* can aid this problem, and can also provide a pivotal change to the scheme when applying it to high dimensional systems.

3.1.1. Localisation. Localisation, a scheme frequently used in the EnKF for high dimensional systems, can also be applied to the ETPF [5]. In the simplest form, localisation applied to the ETPF means that one can reduce the computational cost of designing a multivariate coupling to d times the cost of designing a univariate coupling. Localisation allows one to construct an individual transformation in (3.7) for each of the d components of a multivariate X_{L,t_k} . A simple definition for the localisation matrix C [5] that describes the spatial correlation structure of the ensemble $\{X_{L,t_k}\}_{i=1,\dots,N}$ could be

$$(3.14) \quad C_{m,n} = \begin{cases} 1 - \frac{1}{2} \left(\frac{s_{m,n}}{r_{loc,c}} \right), & \left(\frac{s_{m,n}}{r_{loc,c}} \right) \leq 2, \\ 0, & \text{otherwise.} \end{cases}$$

Here $m, n = 1, \dots, d$ are the indices of the spatial components of X_{L,t_k} ,

$$(3.15) \quad s_{m,n} = \min \{ |m - n - N|, |m - n|, |m - n + N| \},$$

and $r_{loc,c}$ is a constant. The above form for $C_{m,n}$ explicitly takes spatial-periodicity into account. One can now decompose the linear transport problem in (3.13) into d separate linear transport problems, to find a coupling matrix $T_{i,j}(m)$, $i = 1, \dots, N$, $j = 1, \dots, N$, for each component $m = 1, \dots, d$. The objective of these linear transport problems is minimising the cost function

$$(3.16) \quad \sum_{i=1}^N \sum_{j=1}^N T_{i,j}(m) \left\| X_{L,t_k}^i - X_{L,t_k}^j \right\|_m^2,$$

where

$$(3.17) \quad \left\| X_{L,t_k}^i - X_{L,t_k}^j \right\|_m^2 = \sum_{n=1}^d C_{m,n} (X_{L,t_k}^i(n) - X_{L,t_k}^j(n))^2,$$

subject to the constraints

$$(3.18) \quad \sum_{i=1}^N T_{i,j}(m) = \frac{1}{N}, \quad \sum_{j=1}^N T_{i,j}(m) = \tilde{w}_L(X_{L,t_k}^i).$$

Here, $X_{L,t_k}^i(m)$ is the m 'th component of X_{L,t_k}^i . Then, one can define the approximation of the marginal posterior mean for each $m = 1, \dots, N$ as

$$(3.19) \quad \bar{X}_{L,t_k}(m) = \frac{1}{N} \sum_{j=1}^N \tilde{X}_{L,t_k}^j(m),$$

where the transformed components are given by

$$(3.20) \quad \tilde{X}_{L,t_k}^j(m) = \sum_{i=1}^N P_{i,j}(m) X_{L,t_k}^i(m),$$

and $P_{i,j}(m) = T_{i,j}(m)N$. Note the cost functions in (3.17) do not achieve the minimum of (3.13). When $r_{loc,c} = 0$, exhibiting the most computationally efficient scenario, one has the interesting case where d univariate linear transport problems need to be solved, thus transforming all components individually. One can simply use the univariate algorithm in [20], mentioned in the last section, with a computational cost of $O(N \log(N))$, for each linear transport problem, to get an overall $O(dN \log(N))$ computational cost. In practice, when $r_{loc,c} = 0$, one can also reorder each of the transformed sets of components into the rank structure of the original ensemble. This preserves the copula structure [21] of the original ensemble.

If the model costs of a system are less than that of the multivariate optimal transportation, using $r_{loc,c} = 0$ is the only case in which the model costs can return to being the dominative cost in the ETPF estimator. Despite this, it is very reasonable to imagine

that this model cost will dominate that of the optimal transportation in some systems, especially for high dimensional Partial Differential Equations (i.e. where γ is high). Localisation does have an effect on the performance of the ETPF by adding bias into the posterior mean approximation stemming from the fact that one is generating deterministic couplings $T_{i,j}$ that are minimising different (simplified) cost functions to the full, multivariate one in (3.13). This bias is thus caused by the decay in correlation between the components. Despite this, numerical experiments conducted in [5] find the localised ETPF to be effective even in the chaotic, highly nonlinear Lorenz Equations.

Localisation is also needed in the likelihood evaluation of multivariate particles in the ETPF. Although this is not critical to the aim of this paper, only briefly covered here, it is essential for the ETPF to be able to successfully filter high dimensional systems due to the curse of dimensionality. Standard Sequential Monte Carlo (SMC) methods fail to track high dimensional systems due to exponentially degenerate importance weights. However, while there have only been some suggestions for a solution to this problem in SMCs, such as in [18], one can alter the above localisation scheme in the ETPF to solve this problem swiftly [5]. It is also needed to reduce the computational cost of likelihood evaluations when the dimension of the state space is greater than the sample size ($d > N$). For each component ($m = 1, \dots, d$) in the particles $\{X_{L,t_k}^i\}_{i=1,\dots,N}$, generate an separate importance weight given by

$$(3.21) \quad w_L(X_{L,t_k}^i(m)) \propto \frac{1}{\sqrt{2\pi|R|^y/2}} e^{-\frac{1}{2} \left(H(X_{L,t_k}^i) - Y_{t_k} \right)^T (\tilde{C}_m) R^{-1} \left(H(X_{L,t_k}^i) - Y_{t_k} \right)},$$

where

$$(3.22) \quad (\tilde{C}_m)_{n,n} = \begin{cases} 1 - \frac{1}{2} \left(\frac{s_{m,n}}{r_{loc,R}} \right), & \left(\frac{s_{m,n}}{r_{loc,R}} \right) \leq 2, \\ 0, & \text{otherwise,} \end{cases}$$

for $n = 1, \dots, N$ (\tilde{C}_m is diagonal) and the value of $r_{loc,R}$ can be independent to $r_{loc,c}$. Of course, H should be a local operator, see [1] for details of the use of localisation within the EnKF. These weights are then used in the constraints in the linear transport problems for each individual component transformation in (3.18). The two ‘radii’ of localisation, $r_{loc,c}$ and $r_{loc,R}$ will henceforth be referred to as the particular settings of localisation used.

4. Multilevel Ensemble Transform Particle Filter (MLETPF). The proposed multilevel ETPF framework is demonstrated in this section. It creates an estimator consistent with the standard ETPF estimator in (3.6), for the same discretization accuracy level, L . The term ‘single level’ estimator will henceforth be a reference to the corresponding standard ETPF estimator, conditioned on the same observations, with the same discretization level L and variance as the proposed MLETPF estimator. The general premise of the MLETPF is to run $L + 1$ independent ETPF estimators, with N_l samples, forward in time (and space), in the coupled multilevel framework. When updating the weights of each particle in each estimator, the same method as the ETPF

holds for each of the $L + 1$ estimators. Thus, we define the MLETPF estimator of $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, as the following, where the importance weights $\tilde{w}_l(X_{l,t_k})$ target η_{l,t_k} , the posterior of each d -dimensional discretization X_{l,t_k} , given the observations, Y_{t_1, \dots, t_k} , $k \in [1, N_y]$,

$$(4.1) \quad \bar{X}_{L,t_k} = \left(\sum_{i=1}^{N_0} \tilde{w}_0(X_{0,t_k}^i) X_{0,t_k}^i \right) + \sum_{l=1}^L \left(\sum_{i=1}^{N_l} \left(\tilde{w}_l(X_{l,t_k}^i) X_{l,t_k}^i - \tilde{w}_{l-1}(X_{l-1,t_k}^i) X_{l-1,t_k}^i \right) \right)$$

We assume here that $h_0 \leq \Delta t$, where $\Delta t = t_{k+1} - t_k$ for all $k \in [1, N_y - 1]$, so that all of the $L + 1$ estimators are conditioned on the same observations. This does mean that one has to set a bound on the frequency of the data assimilation, given the time-step of the minimum level, h_0 . We note that it is possible to adjust the framework here slightly to incorporate frequent observations only available on finer levels at certain times. This could be done by only updating the weights for finer ensembles at those observations and then proceeding with the ensemble transform stages when both the coarse and fine levels in each difference estimator have had an importance weight update.

One notes that as the standard ETPF estimator for each level of discretization, $l \geq 0$, is consistent with $\mathbb{E}_{\eta_{l,t_k}}[X_{l,t_k}]$, the above estimator is consistent with the $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, given the linearity of expectation shown in (2.2). Here, each of the particles from the fine and coarse ensembles in each of the multilevel difference estimators are positively correlated in between assimilation steps as in the standard MLMC method. This correlation is required for the variance of each difference estimator to decay with $l \rightarrow \infty$ as discussed in the opening section. However, now in the ETPF context, when one comes to transform the fine and coarse ensembles in each multilevel difference estimator, the two ensembles cannot be transformed independently of one another, and need to have a positive correlation imparted between them ready for the next phase of particle propagation, especially if the transformations are happening frequently. If the random input to the system is simply a random initial condition, in a system with no stochastic forcing, these particles from the fine and coarse ensembles will certainly diverge instantly if they are not positively correlated after the ensemble transformations. In this paper, this positive correlation is achieved using a multilevel coupling step after the standard ensemble transform stage. This requires one to first carry out the ensemble transform (3.7) on the coarse and fine ensembles, $\{X_{l-1,t_k}^i\}_{i=1, \dots, N_l}$, $\{X_{l,t_k}^i\}_{i=1, \dots, N_l}$ with weights $\{w_{l-1}(X_{l-1,t_k}^i)\}_{i=1, \dots, N_l}$, $\{w_l(X_{l,t_k}^i)\}_{i=1, \dots, N_l}$ respectively, to get evenly weighted particles, $\{\tilde{X}_{l-1,t_k}^i\}_{i=1, \dots, N_l}$ and $\{\tilde{X}_{l,t_k}^i\}_{i=1, \dots, N_l}$. If localisation is needed, one can implement this with the required parameters on both ensemble transforms as in the last section. It is very important that the same localisation settings are used on all estimators, so that the overall MLETPF estimator is consistent with the single level ETPF estimator with the same localisation settings. The key point of this being that the fine and coarse ensembles from each discretization level will have the same systematic localisation bias as one another. This means, such as with the discretization bias, that the localisation biases can cancel each other out in the telescoping sum of estimators (4.1), leaving only the systematic localisa-

tion bias of the finest level equal to that of the localised single level estimator. At this point, one notes that (4.1) becomes

$$(4.2) \quad \bar{X}_{L,t_k} = \left(\frac{1}{N_0} \sum_{i=1}^{N_0} \tilde{X}_{0,t_k}^i \right) + \sum_{l=1}^L \left(\frac{1}{N_l} \sum_{i=1}^{N_l} (\tilde{X}_{l,t_k}^i - \tilde{X}_{l-1,t_k}^i) \right).$$

Now one needs to positively couple the fine and coarse ensembles of transformed particles from each estimator above. We propose to build another coupling between \tilde{X}_{l,t_k} and \tilde{X}_{l-1,t_k} , denoted by $T_{i,j}^{F/C}$, that minimises the cost function

$$(4.3) \quad \sum_{i=1}^{N_l} \sum_{j=1}^{N_l} T_{i,j}^{F/C} \left\| \tilde{X}_{l,t_k}^i - \tilde{X}_{l-1,t_k}^j \right\|^2,$$

with constraints

$$(4.4) \quad \sum_{i=1}^N T_{i,j}^{F/C} = \frac{1}{N_l}, \quad \sum_{j=1}^N T_{i,j}^{F/C} = \frac{1}{N_l}.$$

This is an assignment problem and in the multivariate case it can be solved by the Hungarian algorithm [15] with a computational cost equal to the multivariate linear transport problem algorithms discussed previously and so is the same order of magnitude as the corresponding ensemble transform stage in the standard ETPF method. In the univariate case, one can simply use the cheap algorithm in [20], exactly like the ensemble transform stage. One notes that the above assignment problem returns a coupling with one element in each row and column ($\frac{1}{N_l}$), resulting in particles simply being reordered and not transformed. This therefore returns exactly the same transformed particles in each ensemble. The reordering can be seen as finding the transformation matrix $P_{i,j}^{F/C} = T_{i,j}^{F/C} N_l$ and then applying the standard ensemble transform, in (3.7), to both the fine and coarse transformed ensembles to get new ensembles of $\{\tilde{X}_{l,t_k}^i\}_{i=1,\dots,N_l}$ and $\{\tilde{X}_{l-1,t_k}^i\}_{i=1,\dots,N_l}$ which are now positively correlated. Each multi-level difference estimator can now be estimated by

$$(4.5) \quad \hat{X}_{l,t_k} = \frac{1}{N_l} \sum_{j=1}^{N_l} (\tilde{X}_{l,t_k}^j - \tilde{X}_{l-1,t_k}^j).$$

Using a calculation similar to (3.12), we now show that the estimator in (4.5) is consistent with the term

$$(4.6) \quad \left(\sum_{i=1}^{N_l} (\tilde{w}_l(X_{l,t_k}^i) X_{l,t_k}^i - \tilde{w}_{l-1}(X_{l-1,t_k}^i) X_{l-1,t_k}^i) \right)$$

from Equation (4.1). Let $T_{i,j}^F$ and $T_{i,j}^C$ ($i, j = 1, \dots, N$) be the coupling matrices used for the ensemble transform on the finer and coarse ensembles respectively, then

$$(4.7) \quad \begin{aligned} \hat{X}_{l,t_k} &= \frac{1}{N_l} \sum_{j=1}^{N_l} (\tilde{X}_{l,t_k}^j - \tilde{X}_{l-1,t_k}^j) = \sum_{j=1}^{N_l} \sum_{i=1}^{N_l} T_{i,j}^{F/C} (\tilde{X}_{l,t_k}^i - \tilde{X}_{l-1,t_k}^i) \\ &= \frac{1}{N_l} \sum_{i=1}^{N_l} \sum_{j=1}^{N_l} N_l (T_{i,j}^F X_{l,t_k}^j - T_{i,j}^C X_{l-1,t_k}^j) = \sum_{j=1}^{N_l} (\tilde{w}_l(X_{l,t_k}^j) X_{l,t_k}^j - \tilde{w}_{l-1}(X_{l-1,t_k}^j) X_{l-1,t_k}^j). \end{aligned}$$

The estimator in (4.2) can therefore be written as

$$(4.8) \quad \bar{X}_{L,t_k} = \left(\frac{1}{N_0} \sum_{i=1}^{N_0} \tilde{X}_{0,t_k}^i \right) + \sum_{l=1}^L \left(\frac{1}{N_l} \sum_{i=1}^{N_l} (\tilde{X}_{l,t_k}^i - \tilde{X}_{l-1,t_k}^i) \right),$$

with covariance

$$(4.9) \quad \text{Cov}[\bar{X}_{L,t_k}] = \sum_{l=0}^L \frac{\mathbb{V}_l}{N_l},$$

where

$$(4.10) \quad \mathbb{V}_l = \begin{cases} \text{Cov}[\tilde{X}_{0,t_k}], & l = 0, \\ \text{Cov}[\tilde{X}_{l,t_k} - \tilde{X}_{l-1,t_k}], & l > 0, \end{cases}$$

due to the fact that each estimator in (4.8) is independent. The above estimator is consistent with the single level (localised, with the same settings as used in the MLETPF) ETPF estimator due to (4.7) and (4.1). Therefore, in the absence of localisation, it is also consistent with $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$.

The multilevel coupling $T_{i,j}^{F/C}$ minimises the expected distance between the two transformed ensembles, which maximises the covariance between them via (3.9) and then finally minimises V_l . The multilevel coupling procedure above minimises \mathbb{V}_l in each multilevel difference estimator; we also ensure that pairs of particles in the coarse and fine ensembles are positively coupled in between assimilation steps (by using the same random input). The aim of this is to make the covariance \mathbb{V}_l decrease at an asymptotic rate $O(h_l^\beta)$ ($\beta > 0$) required for the variance reduction of the multilevel framework to work. This is because the fine and coarse ensembles are coupled both in between assimilation steps and during them via the coupling. This will be demonstrated in numerical experiments later in the paper. Designing this coupling between the both transformed ensembles \tilde{X}_{l-1,t_k} and \tilde{X}_{l,t_k} is the key to the proposed MLETPF method, and enforces correlation amongst both transformed ensembles whilst remaining consistent with the single level ETPF estimator. Most importantly it also suits the ETPF method since the coupling can be generated simply and cheaply when using the $r_{loc,c} = 0$ localisation that can be used freely in the ETPF. This will be shown later.

4.1. Algorithm. In this section, we present an algorithm to implement the MLETPF in practice. In this paper, a pre-defined recurrence relation for the decay of N_l as $l \rightarrow \infty$ will be set ($N_{l+1} = f(N_l)$) and these sample sizes will be kept fixed throughout the filtering process. The finest discretization level $L > 0$ will be kept arbitrary for now. The algorithm is now presented.

1. Start at t_0 (thus $k = 0$) and with $l = 0$. Choose N_0 .
2. Calculate $N_l = \lceil f(N_{l-1}) \rceil$ if $l > 0$. Sample $\{X_{0,t_0}^i\}_{i=1,\dots,N_0} \sim \pi^0$ if $l = 0$ or $\{X_{l,t_0}^i, X_{l-1,t_0}^i\}_{i=1,\dots,N_l} \sim \pi^0$, where $X_{l,t_0}^i = X_{l-1,t_0}^i$ if $l > 0$, at time t_0 .
3. Propagate all samples forward according to system dynamics until time t_{k+1} . If $l > 0$, the fine and coarse pairs of samples in each estimator must be coupled by using the same random input.
4. Derive the normalised importance weights for $X_{l,t_{k+1}}$:

$$(4.11) \quad \tilde{w}_l(X_{l,t_{k+1}}^i) = \frac{w_l(X_{l,t_{k+1}}^i)}{\sum_{j=1}^{N_l} w_l(X_{l,t_{k+1}}^j)}$$

5. Transform the ensembles, $\{X_{l,t_k}^i, \tilde{w}_l(X_{l,t_k}^i)\}_{i=1,\dots,N_l}$ and $\{X_{l-1,t_k}^i, \tilde{w}_{l-1}(X_{l-1,t_k}^i)\}_{i=1,\dots,N_l}$ into the evenly weighted ensembles, $\{\tilde{X}_{l,t_k}^i\}_{i=1,\dots,N_l}$ and $\{\tilde{X}_{l-1,t_k}^i\}_{i=1,\dots,N_l}$, using the linear transformation in (3.7). Couple them using the multilevel coupling matrix $T_{i,j}^{F/C}$ to produce reordered ensembles $\{\tilde{\tilde{X}}_{l,t_k}^i\}_{i=1,\dots,N_l}$ and $\{\tilde{\tilde{X}}_{l-1,t_k}^i\}_{i=1,\dots,N_l}$.
6. Move on to Step 7 if $l = L$ or if not, iterate $l+1$ then repeat steps 2-6 for $k = 0$ or steps 3-6 for $k > 0$.
7. Iterate $k+1$. Start again from step 3 with $l = 0$. The MLETPF approximation of $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$ is given by (4.8).

4.2. Computational Cost of the MLETPF and ETPF. As previously noted, in a multidimensional case, it is computationally expensive to generate the coupling matrices needed to couple the fine and coarse transformed particles. Localisation is used to reduce the computational cost of the ensemble transforms down to $O(dN\log(N))$ when $r_{loc,c} = 0$ in the standard and multilevel ETPF methods, and this too can also reduce the computational cost of generating the multilevel coupling $T_{i,j}^{F/C}$. When localisation is used along with $r_{loc,c} = 0$, one can break the multivariate coupling $T_{i,j}^{F/C}$ down into d separate univariate couplings. As the components are transformed individually, one can simply find a coupling $T_{i,j}^{F/C}(m)$ for each individual component m in the transformed coarse and fine ensembles with the cost function,

$$(4.12) \quad \sum_{i=1}^{N_l} \sum_{j=1}^{N_l} T_{i,j}^{F/C}(m) (\tilde{X}_{l,t_k}^i(m) - \tilde{X}_{l-1,t_k}^j(m))^2,$$

and the same constraints as in (4.4). Each of these couplings can again be found using the cheap, $O(N_l \log(N_l))$ (for each l) univariate algorithm in [20]. One can then reorder / transform each component of fine and coarse ensembles separately using the same methodology as in the last section. This performs a similar role as resampling

N_l particles from $F_{f,m}^{-1}(u)$ and $F_{c,m}^{-1}(u)$ where $F_{f,m}^{-1}$ / $F_{c,m}^{-1}$ are the marginal (empirical) inverse cumulative distribution functions of $\tilde{X}_{l,t_k}(m)$ and $\tilde{X}_{l-1,t_k}(m)$ respectively. Here the same uniform variate $u \in [0, 1]$ is used for each pair of the N_l samples. If localisation is carried out along with $r_{loc,c} > 0$, the computational cost of the optimal transport in the standard ETPF and thus the multilevel coupling, minimizing the full multivariate cost function in (4.3), in the MLETPF will rise to $O(dN_l^3 \log(N_l))$. This is because d different localised, but still multivariate, optimal transportation problems will have to be solved. Therefore, in this scenario, the model costs are likely to be dwarfed by these optimal transportation costs, which are fixed by definition. In this case there is no justification for implementing the multilevel framework as it only aims to reduce model cost and not the optimal transportation computational expense. However to consider the case where model cost does dominate that of the multivariate optimal transportation, the full multivariate coupling will be demonstrated in the numerical experiments at the conclusion of this paper.

This paper now considers the overall computational cost of the ETPF and MLETPF estimators when localised, with $r_{loc,c} = 0$. In this case, as explained above, one can reduce the computational expense of not only the ensemble transform stage, but the multilevel coupling stage as well, in the MLETPF scheme from a potential $O(N_l^3 \log(N_l))$ to $O(dN_l \log(N_l))$ for each multilevel difference estimator. This is enough, with suitable assumptions, to expect that the model computational cost bounds for the standard MLMC method in Theorem 2.1 are of the same order of magnitude to that of the entire MLETPF, including the ensemble transform and coupling stages, as one can simply ‘hide’ the optimal transportation costs behind the particle propagation costs. This follows from the proposition.

PROPOSITION 4.1. *If $\Delta t \geq h_0$ is constant, and one can bound the computational cost of all N_y ensemble transform / multilevel coupling stages (with the last term being the cost associated to the sort prior to the algorithm) of the MLETPF by,*

$$(4.13) \quad C_{ET} \leq \sum_{l=0}^L \left(dc_1 N_l N_y + de_1 N_l \log(N_l) N_y \right)$$

where c_1 is a positive constant, then the total computational cost of the MLETPF is bounded by,

$$(4.14) \quad C_{MLETPF} \leq \sum_{l=0}^L \left(c_2 N_l C_{l,d} t_{N_y} + de_1 N_l \log(N_l) N_y \right)$$

where $C_{l,d}$ is the cost of propagation of one particle on level l (dependent on d), e_1 is a positive constant and c_2 is a positive constant.

Proof. Let the total computational cost of the MLETPF be given by the sum of the model cost and the ensemble transform cost (including the localised likelihood evaluation in (3.21), which scales at $O(N_l)$ due to the sparse diagonal $d \times d$ matrix \tilde{C} , with $r_{loc,R}$ assumed constant ($r_{loc,R} = O(1)$) and $\ll d$)

$$(4.15) \quad C_{MLETPF} = C_{ET} + C_{MODEL}$$

then bounding C_{ET} as in the claim and using the model cost in (2.8),

$$\begin{aligned}
 C_{MLETPF} &\leq \sum_{l=0}^L \left(N_l (dc_1 N_y + C_{l,d} c_3 t_{N_y}) + de_1 N_l \log(N_l) N_y \right) \\
 (4.16) \quad &\leq \sum_{l=0}^L \left(N_l (dc_4 t_{N_y} + C_{l,d} c_3 t_{N_y}) + de_1 N_l \log(N_l) N_y \right) \\
 &\leq \sum_{l=0}^L \left(c_2 N_l C_{l,d} t_{N_y} + de_1 N_l \log(N_l) N_y \right)
 \end{aligned}$$

Here we assume that $C_{l,d}$ will grow at least linearly with d , $c_4 = \frac{c_1}{\Delta t}$ and that c_3 is a positive constant. \square

The last term of each of the expressions above comes from the sorting of the univariate particles in the cheap algorithm. Although this computational cost scales at $O(N_l \log(N_l))$, the constant e_1 is typically very small with respect to the other constants in the cost expression; furthermore this scaling is a worst case estimate, in many cases this would not be reached. Therefore assuming that this part is less than the remainder of the localised ensemble transform cost, one can bound the overall cost of the MLETPF by the model cost, and keep the corresponding reductions outlined in Theorem 2.1. However, even if this sorting cost does dominate the ensemble transform cost, one still expects to recover computational cost reductions relative to the single level estimator, for a fixed error. Despite the model cost in this case not necessarily being the highest computational expense in the localised MLETPF, with the cost of the sorting / ensemble transform dominating, it is important to note that for a fixed error bound, this sorting cost / ensemble transform will still typically be less than the model computational cost of the single level ETPF. The reductions in Theorem 2.1 would then be a slight underestimation however there would still be evident reductions of computational cost relative to the single level ETPF. This is indeed the case for the numerical examples in the next section, as we show there.

4.3. Numerical Examples. Numerical examples of the MLETPF method, applied to classical data assimilation problems, are given in this section. Three problems will be studied: the multivariate, chaotic Stochastic Lorenz-63 Equations, the univariate, but nonlinear double-well OU Process and the high dimensional Stochastic Lorenz-96 Equations. The algorithm above will be used to generate experimental MLETPF estimators (that are compared against the single level ETPF estimators) for the latter two of the three problems above, along with varying levels of pre-defined of accuracy. This pre-defined level of accuracy, $O(\varepsilon)$, will determine L (the finest level / overall numerical discretization bias), the fixed sample sizes for each estimator in the MLETPF method (N_l) and the fixed overall sample size in the corresponding single level ETPF estimator required to achieve the order of magnitude of this error in both estimators. This follows from the standard Monte Carlo approximation error decomposition given by the Central Limit Theorem, as in [9]. One can then compare the computational cost, given as the number of operations, for both the single level and multilevel estimators, which should be in line with Theorem 2.1 given the same

pre-defined error. The error is not bounded exactly due to variations in the variance at each assimilation step, but a proof of concept from a practitioner's viewpoint can be established.

The error in the estimators will be estimated by the time-averaged root mean square error (RMSE), given by

$$(4.17) \quad RMSE = \sqrt{\frac{1}{N_y} \sum_{k=1}^{N_y} \left\| \bar{X}_{L,t_k} - \mathbb{E}_{\eta_{t_k}}[X_{t_k}] \right\|^2},$$

where η_{t_k} is the posterior distribution of X_{t_k} given the observations Y_{t_1, \dots, t_k} . An approximation of $\mathbb{E}_{\eta_{t_k}}[X_{t_k}]$ will be used in the RMSE calculations above, by computing a standard ETPF estimator for it, of which the numerical discretization bias and sample size produce an estimator with an error orders less than any ε used in the following experiments. Where localisation is used in the single level and multilevel estimators for the problems above, the estimators are inconsistent with $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, but crucially consistent with each other as the same localisation settings are used. The ETPF approximation of $\mathbb{E}_{\eta_{t_k}}[X_{t_k}]$ will use the same localisation settings to correctly compare the ETPF and MLETPF estimators like for like.

The recurrence relation for N_l , $l \geq 1$, given N_0 (dependent on ε), used in both numerical experiments will be set to $N_{l+1} = \lceil N_l M^{-3/2} \rceil$, however the optimality of this depends on the relative value of β with respect to γ [9]. This is only optimal for $\beta > \gamma$, $(\beta + \gamma) = 3$; this is the case in the numerical examples below. One notes that for a RMSE of $O(\varepsilon)$, one requires that $N_l = O(\varepsilon^{-2})$, and that $N = O(\varepsilon^{-2})$ for the single level ETPF. Also, the discretization bias of the multilevel and single level ETPF estimators, from $\mathbb{E}_{\eta_{t_k}}[X_{t_k}]$, should be $O(\varepsilon)$. Thus, for a numerical scheme that has a global discretization bias of $O(t_{N_y} h_L^\alpha)$, one requires

$$(4.18) \quad L = \left\lceil \frac{\log((t_{N_y}^\alpha d)/\varepsilon)}{\alpha \log(M)} \right\rceil$$

for the sum of the multilevel and single level ETPF estimator components bias to be of $O(\varepsilon)$. One could also use the maximum among all components' biases to be a suitable measure here. The sample covariances of the independent ETPF estimators will also be measured by a sum among all components: $Tr(\mathbb{V}_l)$, the trace of the covariance matrix.

For the single level ETPF, using the above analysis, one can see how requiring $h_L = O(\varepsilon)$ and $N = O(\varepsilon^{-2})$, the model cost in Equation (2.9) will be $O(\varepsilon^{-3})$, dominating the localised (with $r_{loc,c} = 0$) ensemble transform cost $O(dN \log(N))$, which is equivalent to $O(d\varepsilon^{-2} \log(\varepsilon^{-2}))$; this supports the point made in the last section. This is also greater than the localised ensemble transform cost of the MLETPF, $O(dN_l \log(N_l))$, again equivalent to $O(d\varepsilon^{-2} \log(\varepsilon^{-2}))$. Thus the computational cost reductions of the MLETPF, even if the sorting cost of the localised ensemble transform dominates, is apparent in these cases. The computational cost for the these

numerical examples is defined as the theoretical number of operations needed to compute each approximation, including the ensemble transform stage and the multilevel coupling for the MLETPF. This is computed simply by inserting a step-counter into the numerical implementation (functions etc.) in the Python code used for these experiments. Finally, $M = 2$ is used for each numerical example.

4.3.1. Stochastic Lorenz-63 Equations. This simple 3-component chaotic nonlinear system in $X = (x, y, z)$,

$$(4.19) \quad \frac{dX}{dt} = \begin{cases} \sigma(y - x) + \phi \frac{dW}{dt}, \\ x(\rho - z) - y + \phi \frac{dW}{dt}, \\ xy - \beta z + \phi \frac{dW}{dt}, \end{cases}$$

with $\rho = 28$, $\sigma = 10$, $\beta = 8/3$, $\phi = 0.4$ will be used to demonstrate the effect of the multivariate multilevel coupling, $T_{i,j}^{F/C}$, without localisation and thus with cost function in (4.3). The localised coupling will also be used as a comparison. Here, the Brownian motion W will be the same for each component to keep the strong nonlinearity in the equations. Computational cost against accuracy comparisons with the standard ETPF method will not be investigated here given the low model cost of this test problem and thus the dominating effects of the multilevel coupling and / or the ensemble transform stage in both the ETPF and the MLETPF will make the model cost reductions of the multilevel framework unnoticeable. The MLETPF estimator with $N_0 = 500$, and $L = 5$, using the full multivariate cost function in (4.3) to find the multilevel couplings $T_{i,j}^{F/C}$ in the aforementioned algorithm, is used to compute an approximation to $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, with X as above, for $k \in [1, 5120]$, with $h_0 = M^{-7} = \Delta t$, and thus $t_{N_y} = 40$. Here, X_L is the solution to the above Lorenz-63 equations using the forward Euler numerical scheme. The reason that we choose the minimum level to be equivalent to $l = 7$ is for stability when using the Euler method. Using different numerical methods, with greater stability at greater time-steps, and thus lower levels would be able to decrease this minimum level. The observations are given by a measurement error with $R = 2I$, where I is the 3×3 identity matrix and weights are thus based on observations of all components x , y and z . Figure 1 shows the mean estimates of $Tr(\mathbb{V}_l)$ ($l \in [1, 5]$) over all assimilation steps $k \in [1, N_y]$. The asymptotic decay of the above estimates show importantly that the multilevel coupling, with the multivariate cost function, successfully produces the variance decay, $Tr(\mathbb{V}_l) = O(h_l^\beta)$, in this case $\beta \approx 1$. The figure also shows the value of β (variance decay) for the case where $r_{loc,c} = 0$ localisation for the coupling. Localisation being used in a problem such as the strongly nonlinear Lorenz 63 system is dangerous due to the decay in correlations between components, but with the parameters above, it is used simply to compare the rates of variance decay with the non-localised case. One recovers $\beta \approx 2$ in the localised case; this showing that refining the optimal transport down to the one dimensional localisation case is beneficial for variance reduction but comes with the sacrifice of inconsistency of the reference “single level” estimator .

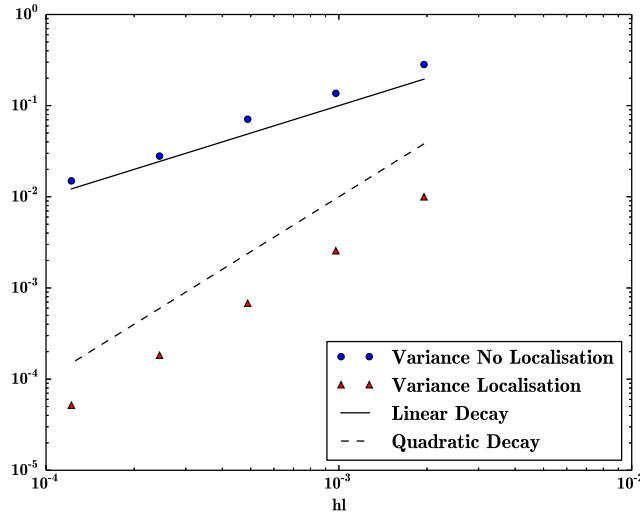


Fig. 1: Mean, over all assimilation steps, $k \in [1, N_y]$, of the estimates of $\text{Tr}(\mathbb{V}_l)$, with $l \in [1, 5]$ for the Stochastic Lorenz-63 Equations. Both the non-localised and the $r_{loc,c} = 0$ localised cases are shown.

4.3.2. Double-well OU Process. This is a univariate test problem that demonstrates the cost effective, consistent, MLETPF estimator of $\mathbb{E}_{\eta_{L,t_k}}[X_{L,t_k}]$, where X_{L,t_k} is a numerical discretized solution to the double-well, nonlinear OU process,

$$(4.20) \quad dX_{t_k} = -V'(X_{t_k})dt + \xi dW_{t_k}$$

$k \in [1, N_y]$ and W_{t_k} is a standard Brownian Motion. Here, $V(X_{t_k}) = \frac{1}{4}X_{t_k}^4 - \frac{1}{2}X_{t_k}^2$. This example uses $h_l = 2^{-4-l}$ but this is arbitrarily chosen. The stochastic forcing is set to $\xi = 0.5$. The observations and assimilation times were given by $R = 0.6$, $t_{N_y} = 50$ where $\Delta t = h_0$ and so $N_y = 800$. The numerical discretizations of X_{t_k} are computed by the Euler-Maruyama numerical scheme. The parameters above produce a stable numerical solution for a single realisation of the above system when using this scheme for the time frame above. A very accurate simulation ($N_0 = 10000$, $L = 7$) of the MLETPF estimator is run to demonstrate the mean asymptotic decay of V_l and $|\hat{X}_{l,t_k}|$ ($l \in [1, 7]$) over all assimilation steps. These are shown in Figure 2. The values of $\alpha \approx 1$, $\beta \approx 2$, are as expected given the Euler-Maruyama global discretization bias of $O(h_l)$ and the additive noise in the OU Process contributing to the variance.

Figure 3 shows the computational cost against the accuracy (RMSE) for the MLETPF and the single level ETPF estimators over varying values of ε . Here one sets $N_0 = \varepsilon^{-2}$ for the MLETPF and $N = \varepsilon^{-2}$ for the single level ETPF estimator. One can clearly see the expected orders of growth for the computational cost of the standard ETPF ($O(\varepsilon^{-3})$, as $\gamma = 1$ and $\alpha = 1$) and the MLETPF ($O(\varepsilon^{-2})$, given that $\gamma < \beta$) that were shown in Theorem 2.1 for the pre-defined RMSE of $O(\varepsilon)$.

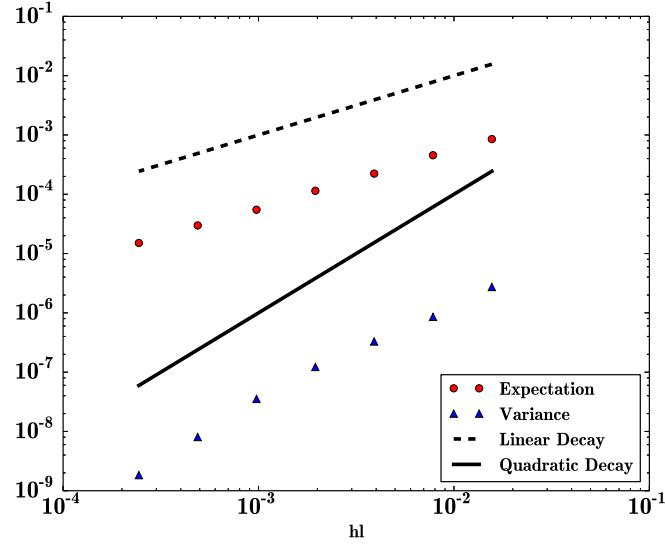


Fig. 2: Mean, over all assimilation steps, $k \in [1, N_y]$, of the estimates of \mathbb{V}_l (variance) and $|\hat{X}_{l,t_k}|$ (expectation) with $l \in [1, 7]$ for the double-well OU process.

4.3.3. The Stochastic Lorenz-96 Equations. The final numerical test for the MLETPF method in this paper is the high dimensional ($d = 40$ in this case) Stochastic Lorenz-96 Equations, given by,

$$(4.21) \quad \frac{dX_j}{dt} = -\frac{(X(j-1)X(j+1) - X(j-2)X(j-1))}{3\Delta X} - X(j) + F + \sigma^2 \frac{dW_j}{dt},$$

with $j \in [1, N_x]$, $N_x = 40$, $N_x \Delta x = 10$ and thus $\Delta X = 0.25$; dW_j/dt are 40 i.i.d. Brownian motions. Here, F is a constant forcing ($F = 8$) and $\sigma^2 = 0.4$. Periodic boundary conditions are used, so that $X(-1) = X(N_x)$. The observations were given by a measurement error of $R = 6I$, where I is the 40×40 identity matrix and assimilation times were set to $t_{N_y} = 100$ where $\Delta t = h_0$ meaning $N_y = 1600$ as $h_l = 2^{-4-l}$ (again simply arbitrary). The Euler-Maruyama method is used to find X_{l,t_k} once again here. In this numerical example, the ETPF and MLETPF estimators use $r_{loc,R} = 1$. The localisation setting of $r_{loc,c} = 0$ is used for both the multilevel and single level ETPF estimators here, due to the model cost, $C_{l,d}$, being simply equal to $h_l^{-1}d$, and thus much lower than that of high optimal transportation costs in multiple dimensions.

Once again, a very accurate simulation ($N_0 = 1000$, $L = 10$) of the MLETPF was generated to demonstrate the mean asymptotic decays of $\sum_{i=1}^{40} |\hat{X}_{l,t_k}(i)|$ and $\text{Tr}(\mathbb{V}_l)$ ($l \in [1, 10]$) over all assimilation steps and these are shown in Figure 4. These follow the same, expected, values of $\alpha \approx 1$, $\beta \approx 2$, as with the last example.

Next, the stability of the MLETPF is considered. Since N_l is fixed and does not change to bound error over time, we can look at the errors from the reference solution, X'_{t_k} , compared to the observational errors, to study the stability of the MLETPF

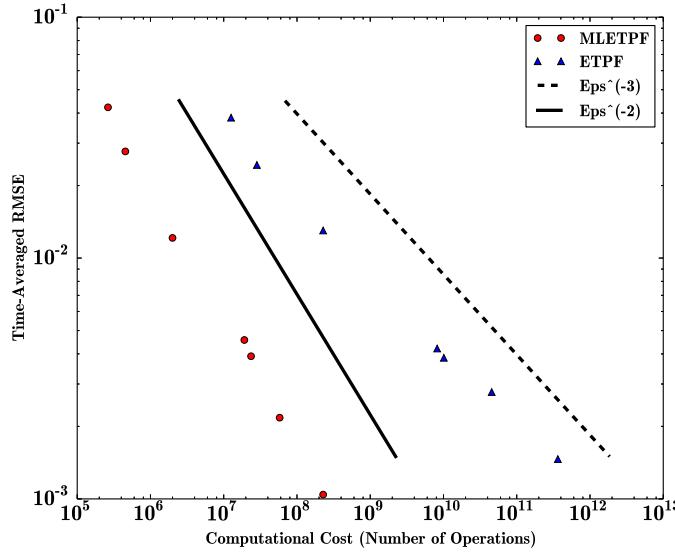


Fig. 3: Computational Cost (number of operations) against the time-averaged RMSE of the ETPF and MLETPF estimators for the double-well OU process. Reference lines show the orders of decay of RMSE^{-2} and RMSE^{-3} .

estimator over time and check that the errors do not increase. One expects that, for a successful particle filter, the errors from the estimator should be less than the observational errors and remain stable. Figure 5 shows this expected behaviour, where the cumulative time-averaged RMSE values from X'_{t_k} , of both the observations and the MLETPF estimator (using arbitrary values of $N_0 = 1000$ and $L = 10$) are shown. For $k \in [1, N_y]$, these are defined to be,

$$(4.22) \quad \sqrt{\frac{1}{k} \sum_{i=1}^k \|Y_{t_i} - X'_{t_i}\|^2}$$

for the observations and,

$$(4.23) \quad \sqrt{\frac{1}{k} \sum_{i=1}^k \|\bar{X}_{L,t_i} - X'_{t_i}\|^2}$$

for the MLETPF estimator. To demonstrate the stability of the variance of the particle filter, cumulative time-averaged RMSE values for the second moments of \bar{X}_{L,t_k} and Y_{t_k} are shown in Figure 6. Finally, to compare the MLETPF with its standard counterpart for this set of equations, Figure 7 shows the computational cost against the accuracy (RMSE) for the standard ETPF and the MLETPF estimators over varying values of ε . Once again, one sets $N_0 = \varepsilon^{-2}$ for the MLETPF and $N = \varepsilon^{-2}$ for the single level ETPF estimator. This follows the successful cost reductions achieved in the last example, defined in Theorem 2.1.

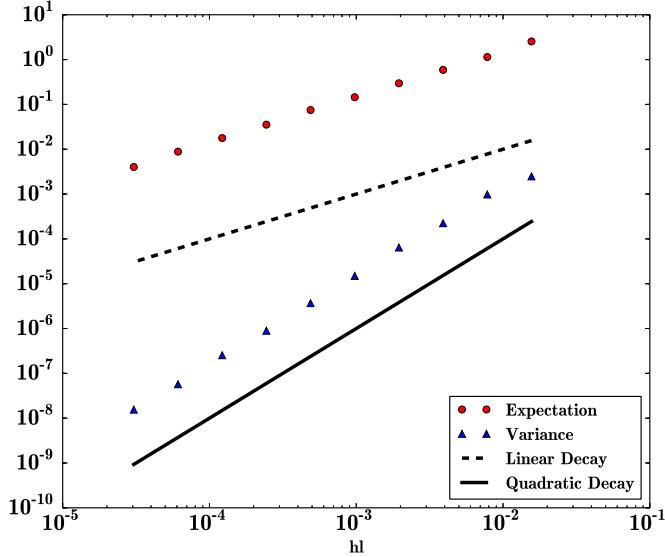


Fig. 4: Mean, over all assimilation steps, $k \in [1, N_y]$, of the estimates of $\text{Tr}(\mathbb{V}_l)$ (variance) and $\sum_{i=1}^{40} |\hat{X}_{l,t_k}(i)|$ (expectation) with $l \in [1, 10]$ for the Stochastic Lorenz-96 Equations.

5. Summary and outlook. This paper has demonstrated a proof of concept for the application of MLMC to nonlinear filtering. The Ensemble Transform Particle Filter (ETPF), coupled with localisation, allows one to simply and cheaply carry out a multilevel coupling between each fine and coarse ensemble in each independent Monte Carlo estimator in the MLMC framework. A recent study has also proposed a framework to apply MLMC to nonlinear filtering with a modified random resampling step in the standard particle filtering methodology to couple particles from coarse and fine levels [12]. In contrast, the coupling in the present paper is designed to minimise the Wasserstein distance between the distributions of these transformed ensembles (in the standard ETPF methodology), originally suggested in [8]. It has been shown through numerical experiments that one can restore positive correlation between fine and coarse ensembles which might have been lost if they had been transformed independently of one another. This in turn satisfies the necessary constraints on the sample variance of each independent multilevel estimator, allowing the proposed MLETPF method to reduce the computational cost of the propagation of particles in the localised ETPF method.

In general, localisation with $r_{loc,c} = 0$ makes the computational cost of this coupling and the ensemble transform in the MLETPF cheap enough that the multilevel framework can return overall computational cost reductions from the standard ETPF methods; the aim of the paper. It must be noted that although this paper has only touched on the case where the very crude $r_{loc,c} = 0$ localisation is considered,

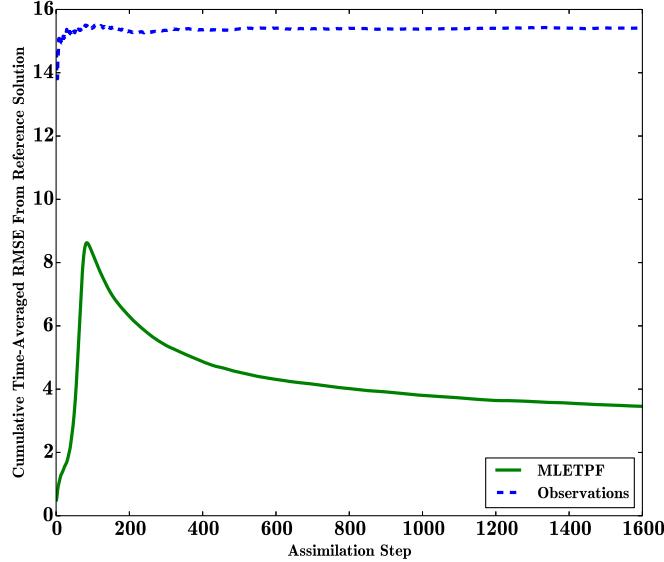


Fig. 5: The cumulative time-averaged RMSE of the observations and MLETPF estimator away from the reference solution, X'_{t_k} , for the Stochastic Lorenz-96 Equations.

due to small model cost test problems, this method could also be applied to high-dimensional systems where the model cost dominates that of the multivariate optimal transportation. One can do this without any crude constraints on $r_{loc,c}$ using the full multivariate coupling methodology presented in this paper and demonstrated numerically. However whether the variance decay of \mathbb{V}_l from such a multivariate coupling would hold, producing as strong results, in the limit of $d \gg N$ is unknown and thus the issue of how one would adjust this coupling to be used alongside other values of $r_{loc,c}$ to reduce the dimensionality of these multivariate couplings remains to be explored.

Iterative and approximate schemes for solving discrete optimal transportation problems have been an area of rapid research in the last few years [17] and this offers the chance to improve the multilevel coupling in the proposed method by reducing computational cost. This could be done by trading-off between the optimality and computational cost of the coupling for each l , e.g. more expensive / optimal couplings for greater l with lower sample sizes N_l .

The form of the coupling used in this paper is simple to implement and has the potential to be used in plenty of applications, in and outside of data assimilation, whenever one wishes to establish consistent correlation between two distributions for variance reduction. Considering an extension for the multidimensional example presented in this paper, one could also apply a spatial multilevel framework, setting the spatial resolution (ΔX_l) to be dependent on the level of discretization, as done in [6, 2] to

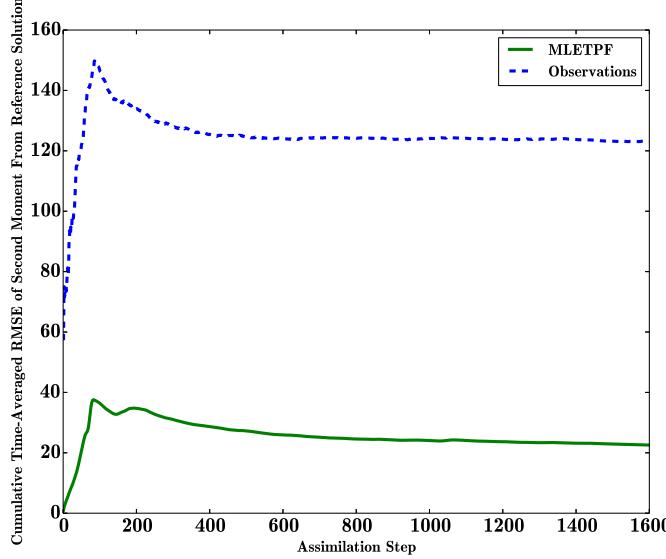


Fig. 6: The cumulative time-averaged RMSE of the second moment of the observations and MLETPF estimator away from the second moment of the reference solution, $(X'_{t_k})^2$, for the Stochastic Lorenz-96 Equations.

gain even more significant cost reductions.

REFERENCES

- [1] J. L. ANDERSON, *Localization and sampling error correction in ensemble kalman filter data assimilation*, Monthly Weather Review, 140 (2012), pp. 2359–2371.
- [2] A. BARTH AND A. LANG, *Multilevel Monte Carlo method with applications to stochastic partial differential equations*, International Journal of Computer Mathematics, 89 (2012), pp. 2479–2498.
- [3] A. BESKOS, A. JASRA, K. LAW, R. TEMPONE, AND Y. ZHOU, *Multilevel Sequential Monte Carlo Samplers*, arXiv preprint arXiv:1503.07259, (2015).
- [4] O. CAPPÉ, S. J. GODSILL, AND E. MOULINES, *An overview of existing methods and recent advances in sequential Monte Carlo*, Proceedings of the IEEE, 95 (2007), pp. 899–924.
- [5] Y. CHENG AND S. REICH, *A McKean optimal transportation perspective on Feynman-Kac formulae with application to data assimilation*, arXiv preprint arXiv:1311.6300, (2013).
- [6] K. A. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients*, Computing and Visualization in Science, 14 (2011), pp. 3–15.
- [7] A. DOUCET AND A. M. JOHANSEN, *A tutorial on particle filtering and smoothing: fifteen years later*, in Oxford Handbook of Nonlinear Filtering, 2011.
- [8] M. GILES, *Multilevel Monte Carlo methods*, in Monte Carlo and Quasi-Monte Carlo Methods, vol. 65, Springer, 2013, pp. 83–103.
- [9] M. B. GILES, *Multilevel Monte Carlo Path Simulation*, Oper. Res., 56 (2008), pp. 607–617.
- [10] M. B. GILES AND B. J. WATERHOUSE, *Multilevel quasi-Monte Carlo path simulation*, Advanced Financial Modelling, Radon Ser. Computat. and Appl. Math., 8 (2009), pp. 165–181.

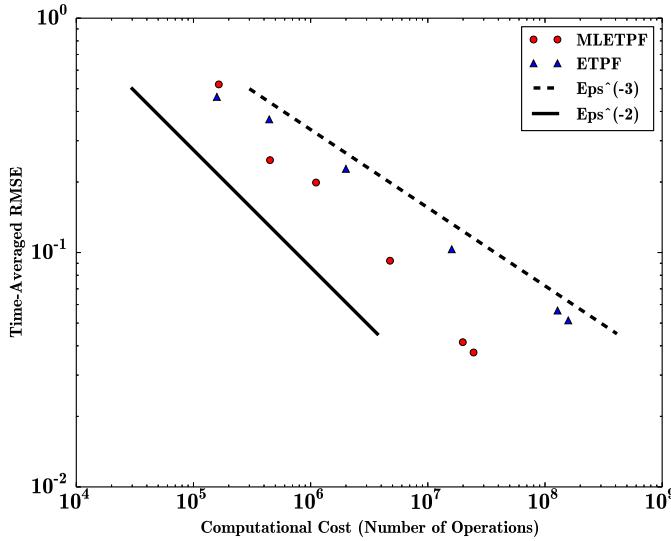


Fig. 7: Computational Cost (number of operations) against the time-averaged RMSE of the ETPF and MLETPF estimators for the Stochastic Lorenz-96 Equations. Reference lines show the orders of decay of RMSE^{-2} and RMSE^{-3} .

- [11] H. HOEL, K. J. H. LAW, AND R. TEMPONE, *Multilevel ensemble Kalman Filtering*, arXiv preprint arXiv:1502.06069, (2015).
- [12] A. JASRA, K. KAMATANI, K. J. H. LAW, AND Y. ZHOU, *Multilevel particle filter*, arXiv preprint arXiv:1510.04977, (2015).
- [13] C. KETELSEN, R. SCHEICHL, AND A. L. TECKENTRUP, *A Hierarchical Multilevel Markov Chain Monte Carlo Algorithm with Applications to Uncertainty Quantification in Subsurface Flow*, Center for Applied Scientific Computing, (2013).
- [14] F. LE GLAND, V. MONBET, AND V. TRAN, *Large sample asymptotics for the ensemble kalman filter*, The Oxford Handbook of Nonlinear Filtering, (2011), pp. 598–631.
- [15] J. MUNKRES, *Algorithms for the assignment and transportation problems*, Journal of the Society for Industrial and Applied Mathematics, 5 (1957), pp. 32–38.
- [16] O. PELE AND M. WERMAN, *Fast and robust earth mover's distances*, in Computer Vision, IEEE 12th International Conference, 2009, pp. 460–467.
- [17] R. RAVI, *Iterative Methods in Combinatorial Optimization*, in STACS'12 (29th Symposium on Theoretical Aspects of Computer Science), vol. 14, 2012, pp. 24–24.
- [18] P. REBESCHINI AND R. VAN HANDEL, *Can local particle filters beat the curse of dimensionality?*, arXiv preprint arXiv:1301.6585, (2013).
- [19] S. REICH, *A nonparametric ensemble transform method for Bayesian inference*, SIAM J. Sci. Comput., 35 (2013), pp. A2013–A2024.
- [20] S. REICH AND C. J. COTTER, *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge Univ. Press, 2015.
- [21] R. SCHEFZIK, T. L. THORARINSDOTTIR, AND T. GNEITING, *Uncertainty quantification in complex simulation models using ensemble copula coupling*, Statist. Sci., 28 (2013), pp. 616–640.
- [22] C. VILLANI, *Optimal transport: Old and New*, vol. 338, Springer Science & Business Media, 2008.